arXiv:1608.00191v1 [cs.IT] 31 Jul 2016

New MDS codes with small sub-packetization and near-optimal repair bandwidth

Venkatesan Guruswami and Ankit Singh Rawat

Computer Science Department,
Carnegie Mellon University,
Pittsburgh, 15213.
E-mail: venkatg@cs.cmu.edu, asrawat@andrew.cmu.edu

August 2, 2016

Abstract

An (n, M) vector code $\mathcal{C} \subseteq \mathbb{F}^n$ is a collection of M codewords where n elements (from the field $\mathbb{F}$) in each of the codewords are referred to as code blocks. Assuming that $\mathbb{F} \cong \mathbb{B}^\ell$, the code blocks are treated as ℓ-length vectors over the base field $\mathbb{B}$. Equivalently, the code is said to have the sub-packetization level ℓ. This paper addresses the problem of constructing MDS vector codes which enable exact reconstruction of each code block by downloading small amount of information from the remaining code blocks. The repair bandwidth of a code measures the information flow from the remaining code blocks during the reconstruction of a single code block. This problem naturally arises in the context of distributed storage systems as the node repair problem [4]. Assuming that $M = |\mathbb{B}|^{k\ell}$, the repair bandwidth of an MDS vector code is lower bounded by $\left(\frac{n-1}{n-k}\right) \cdot \ell$ symbols (over the base field $\mathbb{B}$) which is also referred to as the cut-set bound [4]. For all values of n and k, the MDS vector codes that attain the cut-set bound with the sub-packetization level $\ell = (n-k)^{\lceil n/(n-k) \rceil}$ are known in the literature [23, 35].

This paper presents a construction for MDS vector codes which simultaneously ensures both small repair bandwidth and small sub-packetization level. The obtained codes have the smallest possible sub-packetization level $\ell = O(n-k)$ for an MDS vector code and the repair bandwidth which is at most twice the cut-set bound. The paper then generalizes this code construction so that the repair bandwidth of the obtained codes approach the cut-set bound at the cost of increased sub-packetization level. The constructions presented in this paper give MDS vector codes which are linear over the base field $\mathbb{B}$.

1 Introduction

Maximum distance separable (MDS) codes are considered to be an attractive solution for information storage as they operate at the optimal storage vs. reliability trade-off given by the Singleton bound [15]. For a given amount of information to be stored and available storage space, the MDS codes can tolerate the maximum number of worst case failures without losing the stored information. However, the applicability of the MDS codes in modern storage systems also depends on their ability to efficiently regenerate parts of a codeword from the rest of the codeword. Consider a distributed storage system which employs an MDS code to store information over a network of storage nodes such that each storage node stores a part of a codeword from the MDS code. Exact regeneration (repair) of the content stored in a node with the help of the content stored in the remaining nodes is useful to reinstate the system in the event of a permanent node failure. Similarly, this also enables access to the information stored on a temporarily unavailable node with the help of available nodes in the system. Therefore, among all MDS codes, the ones with more efficient exact repair mechanisms are preferred for deployment in modern distributed storage systems.

In [4], Dimakis et al. study the repair problem in distributed storage systems and introduce *repair bandwidth*, the amount of data downloaded during a node repair, as a metric to compare various codes in terms of the efficiency of their node repair mechanisms. Let $\mathcal{C} \subseteq \mathbb{F}^n$ be an MDS code with $|\mathcal{C}| = |\mathbb{F}|^k$ codewords, each of length n (over $\mathbb{F}$). Given a file $\mathbf{f} \in \mathbb{F}^k$, it is mapped to a codeword in $\mathcal{C}$. Subsequently, the n symbols (over $\mathbb{F}$) in the associated codeword are stored in n distinct storage nodes in the system. For an MDS code, it is straightforward to achieve a repair bandwidth of k symbols (over $\mathbb{F}$) by contacting any k remaining nodes and downloading the k distinct symbols stored on these nodes. This follows from the fact that any k symbols of a codeword from an MDS codes are sufficient to reconstruct the *entire* codeword. Note that the repair bandwidth of k symbols (over $\mathbb{F}$) is the best possible if we are allowed to contact only k remaining storage nodes during the repair process. Furthermore, it is not possible to regenerate a code symbol by contacting less than k remaining code symbols of a codeword in an MDS code. This motivates Dimakis et al. to look for potentially lowering the repair bandwidth for repair of a single node by contacting $d \geq k$ remaining nodes in the system and downloading partial data stored on each of the contacted nodes.

Assuming that the MDS code $\mathcal{C}$ is defined over the field $\mathbb{F} \cong \mathbb{B}^\ell$, we can view each of the n symbols (over $\mathbb{F}$) in a codeword as an ℓ-length vector over the base field $\mathbb{B}$. Given this vector representation of the MDS code, the repair bandwidth of an MDS code is lower bounded by [4, 11]

$$\left(\frac{d}{d-k+1}\right) \cdot \ell \quad \text{symbols (over } \mathbb{B}\text{).} \tag{1}$$

In the particular case, when $d = n-1$, i.e., all the remaining nodes in the system are contacted during the repair process, the bound on the repair bandwidth reduces to $\left(\frac{n-1}{n-k}\right) \cdot \ell$ symbols (over $\mathbb{B}$).

The bound in (1) is referred to as the cut-set bound in the literature. The problem of constructing MDS codes with optimal repair-bandwidth (cf. (1)) has been explored in [3, 8, 17, 18, 22, 23, 28, 35] and references therein. Note that as the number of the nodes contacted during the repair process d gets larger, the optimal repair bandwidth defined by the cut-set bound becomes significantly smaller than the naive repair bandwidth of k symbols (over $\mathbb{F}$) or $k\ell$ symbols (over $\mathbb{B}$).

This paper explores a trade-off between the sub-packetization level ℓ and the repair bandwidth for the MDS codes. The MDS codes that work with small sub-packetization level in addition to having small repair bandwidth are of great practical importance in distributed storage systems. The smaller sub-packetization leads to easier system implementation as it provides the system designer with greater flexibility in terms selecting various system parameters. As an example, one does not have to combine the data from multiple different sources to meet the larger sub-packetization requirement in order to be able to enable efficient repair mechanism. Note that for the given system parameters n, k and d, we require $k\ell$ symbols (over $\mathbb{B}$) worth of data to store (using an MDS code with sub-packetization level ℓ) so that we utilize the storage

space in the most efficient manner. As an illustration of another practical advantage of having smaller sub-packetization level, consider a scenario where a MDS code requires a large sub-packetization level, e.g., say $\ell \geq 2^n$. This implies that using storage nodes (disks) with storage capacity of ℓ symbols (over $\mathbb{B}$), one can only design a storage system with at most $\log_2 \ell$ nodes. Therefore, larger sub-packetization level can lead to a reduced design space in terms of various system parameters.

Our contributions. We present a new and simple construction for MDS codes which have small sub-packetization level while allowing for exact repair of all code symbols with near-optimal repair bandwidth. This construction highlights a trade-off between the sub-packetization level and the repair bandwidth for exact repair. The construction is obtained by utilizing the parity-check view of a linear code. Assuming that the desired sub-packetization level is ℓ, we start with a parity-check matrix of a simple MDS code (over $\mathbb{F} \cong \mathbb{B}^{\ell}$) of length n which is obtained by stacking ℓ codewords from ℓ independent MDS codes (over $\mathbb{B}$) of length n. We then carefully replace some of the zero entries of this parity-check matrix with non-zeros elements from $\mathbb{B}$ and obtain a parity check matrix of a new MDS code that has an exact repair mechanism with small repair bandwidth. It follows from the construction that the obtained MDS codes are linear over the field $\mathbb{B}$.

We note that throughout this paper we consider the setting with $d = n - 1$, i.e., all the remaining code blocks contribute to the exact repair of a single code block. We list the parameters and some of the exact-repair related properties of the obtained codes in the following.

- **Codes with repair bandwidth at most twice the cut-set bound.** We first present a family of MDS codes that have sub-packetization level $\ell = n - k$ and the repair bandwidth that is strictly less than $2(n-1)$ symbols (over $\mathbb{B}$). Note that this is twice the cut-set bound (cf. (1)) which takes the value $(n-1)$ symbols (over $\mathbb{B}$) for $\ell = n - k$. We argue that the sub-packetization level $\ell = \Omega(n-k)$ is the smallest that we can hope for an MDS code with the aforementioned guarantee on its repair bandwidth (See Appendix A).

- **Codes with repair bandwidth approaching the cut-set bound.** We generalize the ideas used in the construction with $\ell = n - k$ to obtain the MDS codes that have improved repair bandwidth at the cost of increased sub-packetization level. In particular, for an integer $t \geq 2$, we obtain a family of MDS codes with sub-packetization level $\ell = (n-k)^t$ and the repair bandwidth which is at most $(1 + \frac{1}{t})$ times the value of the cut-set bound.

- **Exact repair using repair-by-transfer schemes** The codes presented in this paper are MDS codes defined over $\mathbb{F} = \mathbb{B}^{\ell}$ which are linear over the base field $\mathbb{B}$. These codes are referred to as linear MDS vector codes or linear MDS vector codes in the literature. For these codes, n code blocks are stored in the form of an ℓ-length vectors (over $\mathbb{B}$) in n distinct nodes. The exact repair of each code block in these codes involves downloading a subset of the symbols from the remaining code blocks. Such repair mechanisms are referred to as the *uncoded repair* or *repair-by-transfer* in the literature. The repair-by-transfer schemes form a sub-class of all possible linear repair schemes where a contacted node can potentially send symbols (over $\mathbb{B}$) which are linear combinations of all ℓ symbols of the code-block stored on this node. We note that repair-by-transfer is desirable over other complicated repair schemes due to its operational simplicity and the minimal computation requirements at the contacted nodes.

Organization. We introduce the necessary background along with a discussion on the related work in Section 2. In Section 3 we define the notion of near-optimal exact-repairable MDS codes and summarize the code parameters achievable by our construction. We present the construction of the MDS codes with $\ell = n - k$ and repair bandwidth at most twice the value of the cut-set bound in Section 4. In Section 5 we present the the general construction that gives MDS codes with their repair bandwidth approaching the

optimal repair bandwidth. We conclude the paper in Section 6 where we comment on the constructions of the codes with general values of d (the number of blocks contributing to the repair process) and discuss other directions for future work.

2 Background and related work

In this section we formally introduce vector codes and the related concepts used in this paper. We then describe the exact repair problem in the context of distributed storage systems and survey the related work.

2.1 Preliminaries

Let $\mathbb{1}_{\{\cdot\}}$ denote the standard indicator function. Given two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{B}^{n\ell}$, we defined the (block) Hamming distance between them as $d_{\mathrm{H}}(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^{n} \mathbb{1}_{\{\mathbf{x}_i \neq \mathbf{y}_i\}}$, where for $i \in [n]$, we have $\mathbf{x}_i = (x_{(i-1)\ell+1}, \ldots, x_{i\ell})$ and $\mathbf{y}_i = (y_{(i-1)\ell+1}, \ldots, y_{i\ell})$. For a finite field $\mathbb{B}$, we say that a set of vectors $\mathcal{C} \subseteq \mathbb{B}^{n\ell}$ is an $(n, M, d_{\min}, \ell)_{|\mathbb{B}|}$ *vector code* (or in short, (n, M) vector code) if we have $|\mathcal{C}| = M$ and $d_{\min} := \min_{\mathbf{x}\neq\mathbf{y}\in\mathcal{C}} d_{\mathrm{H}}(\mathbf{x}, \mathbf{y})$. Given a codeword $\mathbf{c} = (c_1, c_2, \ldots, c_{n\ell}) \in \mathcal{C} \subseteq \mathbb{B}^{n\ell}$, we use

$$\mathbf{c}_i = (c_{(i-1)\ell+1}, c_{(i-1)\ell+2}, \ldots, c_{i\ell}) \in \mathbb{B}^{\ell}$$

to denote the i-th *code block* in the codeword. When the code $\mathcal{C}$ spans a $\mathbb{B}$-linear subspace of dimension $K = \log_{|\mathbb{B}|} M$, we call $\mathcal{C}$ to be a *linear* vector code and refer to it as an $[n, K, d_{\min}, \ell]_{|\mathbb{B}|}$ vector code. An $[n, K = \log_{|\mathbb{B}|} M, d_{\min}, \ell]_{|\mathbb{B}|}$ vector code is said to be a *linear MDS vector code* if we have $K = \log_{|\mathbb{B}|} M = k\ell$ and $d_{\min} = n - k + 1$. Note that an $[n, k\ell, d_{\min}, \ell]_{|\mathbb{B}|}$ vector code can be defined by an $(n-k)\ell \times n\ell$ parity-check matrix

$$\mathbf{H} = \begin{pmatrix} H_{1,1} & H_{1,2} & \cdots & H_{1,n} \\ H_{2,1} & H_{2,2} & \cdots & H_{2,n} \\ \vdots & \vdots & \ddots & \vdots \\ H_{r,1} & H_{r,2} & \cdots & H_{r,n} \end{pmatrix} \in \mathbb{B}^{(n-k)\ell \times n\ell}, \tag{2}$$

where each $H_{i,j}$ is an $\ell \times \ell$ matrix with its entries belonging to the finite field $\mathbb{B}$. For a set $\mathcal{S} = \{i_1, i_2, \ldots, i_{|\mathcal{S}|}\} \subseteq [n]$, we define the $(n-k)\ell \times |\mathcal{S}|\ell$ matrix $\mathbf{H}_{\mathcal{S}}$ as follows.

$$\mathbf{H}_{\mathcal{S}} = \begin{pmatrix} H_{1,i_1} & H_{1,i_2} & \cdots & H_{1,i_{|\mathcal{S}|}} \\ H_{2,i_1} & H_{2,i_2} & \cdots & H_{2,i_{|\mathcal{S}|}} \\ \vdots & \vdots & \ddots & \vdots \\ H_{r,i_1} & H_{r,i_2} & \cdots & H_{r,i_{|\mathcal{S}|}} \end{pmatrix} \in \mathbb{B}^{(n-k)\ell \times |\mathcal{S}|\ell}. \tag{3}$$

Note that the matrix $\mathbf{H}_{\mathcal{S}}$ comprises those coefficients in the linear constraints defined by the parity-check matrix $\mathbf{H}$ that are associated with the code blocks indexed by the set $\mathcal{S} \subseteq [n]$. The parity-check matrix $\mathbf{H}$ defines an MDS vector codes if for every $\mathcal{S} \subseteq [n]$ with $|\mathcal{S}| = n - k$, the $r\ell \times r\ell$ sub-matrix $\mathbf{H}_{\mathcal{S}}$ is full rank.

2.2 Exact-repair problem for MDS vector codes

Let $\mathcal{C} \in \mathbb{B}^{n\ell}$ be a vector code with $K = \log_{|\mathbb{B}|} |\mathcal{C}| = \log_{|\mathbb{B}|} M$. Consider an encoding process which encodes a file $\mathbf{f} \in \mathbb{B}^K$ to a codeword $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_n) \in \mathcal{C} \subseteq \mathbb{B}^{n\ell}$, where for every $i \in [n]$ we have $\mathbf{c}_i \in \mathbb{B}^{\ell}$. We require the encoding process to ensure that the original file $\mathbf{f}$ can be reconstructed from any k out of the n code blocks in the codeword $\mathbf{c}$, i.e., for any $\mathcal{K} \subseteq [n]$ such that $|\mathcal{K}| = k$, $\mathbf{f}$ can be reconstructed from

the code blocks $\{\mathbf{c}_i\}_{i\in\mathcal{K}}$. The exact-repair problem imposes the requirement that for every $i \in [n]$ and $\mathcal{R} \subseteq [n]\backslash\{i\}$ with $|\mathcal{R}| = d$, we have a collection of functions, $\left\{h^{(i)}_{j,\mathcal{R}} : \mathbb{B}^{\ell} \rightarrow \mathbb{B}^{\beta_{j,i}}\right\}_{j\in\mathcal{R}}$ such that $\mathbf{c}_i$ is a function of the symbols in the set $\{h^{(i)}_{j,\mathcal{R}}(\mathbf{c}_j)\}_{j\in\mathcal{R}}$. This implies that for every $i \in [n]$, the code block $\mathbf{c}_i$ can be *exactly repaired* (regenerated) by contacting any d out of $n-1$ remaining code blocks in the codeword $\mathbf{c}$ (say indexed by the set $\mathcal{R} \subseteq [n]\backslash\{i\}$) and downloading at most $\sum_{j\in\mathcal{R}} \beta_{j,i}$ symbols (over $\mathbb{B}$) from the contacted code blocks.

In [4], Dimakis et al. formally study the repair problem for vector codes in the setup described above. They introduce *repair bandwidth*, the total number of symbols downloaded during the repair process, as a measure to characterize the efficiency of the repair process[1]. Assuming that we download the same number of symbols from each of the contacted code blocks, i.e., $\beta_{j,i} = \beta$ (symbols over $\mathbb{B}$) for all $j \in \mathcal{R}$, Dimakis et al. obtain the following cut-set bound on the repair bandwidth of an MDS vector code [4].

$$d\beta \geq \left(\frac{d}{d-k+1}\right) \cdot \ell \text{ (symbols over } \mathbb{B}). \tag{4}$$

Interestingly, the lower bound on the repair bandwidth of an MDS vector code given in (4) continues to hold even when the contacted nodes contribute different number of symbols during the repair process [11], i.e., for every $i \in [n]$, we have

$$\sum_{j\in\mathcal{R}} \beta_{j,i} \geq \left(\frac{d}{d-k+1}\right) \cdot \ell \text{ (symbols over } \mathbb{B}), \ \ \forall\, \mathcal{R} \subseteq [n]\backslash\{i\} \text{ s.t. } |\mathcal{R}| = d. \tag{5}$$

The problem of constructing *exact-repairable MDS vector code*, MDS vector codes that enable exact repair of all code blocks, with optimal repair bandwidth (cf. (4)) has been explored by many researchers. In [18], Rashmi et al. present an explicit construction for exact-repairable MDS vector codes. This construction works with the sub-packetization level $\ell = d-k+1$. However, the construction requires $2k-2 \leq d \leq n-1$, which leads to low information rate, i.e., $\frac{k}{n} \leq \frac{1}{2} + \frac{1}{2n}$. Towards constructing high-rate exact-repairable MDS codes with optimal repair-bandwidth, Cadambe et al. [3] show the existence of such codes when sub-packetization level approaches infinity. Motivated by this result, the problem of designing high-rate exact-repairable MDS codes with finite sub-packetization level and optimal repair bandwidth is explored in [2, 8, 17, 22, 23, 28, 32, 34, 35] and references therein.

The code construction based on Hadamard matrices from [17] requires $d = n-1$ and $n-k = 2$. In [28], Tamo et al. propose the zigzag code construction for $d = n-1$ and every value of $r = n-k$. This construction enables exact repair of only k (systematic) code blocks. Wang et al. [31] generalize the zigzag code construction to enable exact repair of all code symbols. The constructions presented in [17,28,31] work with the sub-packetization level ℓ which is exponential in k. For $d = n-1$ and all values of $r = n-k$, Sasidharan et al. [23] construct exact-repairable MDS vector codes that have optimal repair bandwidth and work with the sub-packetization level $\ell = (n-k)^{\lceil \frac{n}{n-k} \rceil}$. Note that for $r = n-k = \Omega(n)$, this construction provides the codes with the sub-packetization level which is polynomial in n. The construction with $d = n-1$ and the similar sub-packetization levels that enable exact-repair of only k systematic nodes are also presented in [2, 32]. The construction from [23] is generalized to work for all possible values of $k \leq d \leq n-1$ with the sub-packetization level $\ell = (d-k+1)^{\lceil \frac{n}{d-k+1} \rceil}$ in [22].

The MDS codes presented in [2, 17, 28, 31, 32] are obtained by designing a suitable generator matrix for these codes. On the other hand, [22, 23] design the proposed codes by constructing parity check matrices

[1]Dimakis et al. consider a broader repair framework, namely *functional repair* framework [4]. Under functional repair framework, $\tilde{\mathbf{c}}_i \in \mathbb{B}^{\ell}$ which may potentially be different from the code block under repair $\mathbf{c}_i \in \mathbb{B}^{\ell}$ is an acceptable outcome of the repair process as long as it preserves certain properties of the original codeword. For further details, we refer the reader to [4, 5]. Here, we note that the lower bounds obtained for the functional repair problem are also applicable to the exact repair problem considered in this paper.

Table 1: Comparison of the proposed code construction with various existing code constructions for MDS codes that have small repair bandwidth for exact repair. We focus only on the setting with $d = n - 1$.

Code construction	Sub-packetization level	Repair bandwidth	Repair by transfer	Information rate
Rashmi et al., 2011 [18]	$\ell = n - k$	$\left(\frac{n-1}{n-k}\right) \cdot \ell$	No	$2 \leq 2k \leq n+1$
Ye and Barg, 2016 [35]	$\ell = (n-k)^{\lceil \frac{n}{n-k} \rceil}$	$\left(\frac{n-1}{n-k}\right) \cdot \ell$	Yes	$1 \leq k \leq n-1$
This paper (design parameter $t \geq 1$)	$\ell = (n-k)^t$	$\leq (1 + \frac{1}{t}) \cdot \left(\frac{n-1}{n-k}\right) \cdot \ell$	Yes	$1 \leq k \leq n-1$

with certain combinatorial structures. We note that in most of these constructions, certain elements in the generator/parity-check matrices are not explicitly specified. These papers argue the existence of good choices for these elements provided that the field size is large enough. Recently, Ye and Barg [35] have presented a fully explicit construction for MDS codes with $d = n - 1$ and the sub-packetization level $\ell = (n-k)^{\lceil \frac{n}{n-k} \rceil}$ by designing the associated parity-check matrices. This construction is closely related to the construction presented in [23] in terms of the combinatorial structure of the parity-check matrix. We also note that the construction from [35] also works for general values of $k \leq d \leq n-1$ with suitably modified sub-packetization levels similar to the sub-packetization levels used in [22]. In Table 1 we summarize code parameters of state-of-the-art constructions in different settings.

Some converse results on the sub-packetization level that is necessary for an MDS code to attain the cut-set bound are presented in [9, 29]. For $d = n-1$, Goparaju et al. [9] show that an exact-repairable MDS code that downloads the same number of symbols from each of the contacted code blocks and employs linear repair schemes satisfies the following bound on its sub-packetization level.

$$k \leq 2(\log_2 \ell)\left(\left\lfloor \log_{\frac{n-k}{n-k-1}} \ell \right\rfloor + 1\right). \tag{6}$$

On the other hand, Tamo et al. [29] obtain the following lower bound on the sub-packetization level of an MDS vector code which enables exact repair using repair-by-transfer schemes.

$$\ell \geq (n-k)^{\frac{k}{n-k}}. \tag{7}$$

Note that repair-by-transfer schemes constitute a sub-class of all possible linear repair schemes. In light of the bound in (7), the MDS codes obtained in [23, 35] enable repair-by-transfer mechanisms with optimal repair bandwidth and near-optimal sub-packetization level. However, this sub-packetization level can be prohibitively large for some storage systems, especially when the code has high rate or equivalently has small value of $r = n - k$. This motivates us to explore the question of designing MDS codes that work with small sub-packetization level and provide repair-by-transfer mechanism for exact repair problem without incurring much degradation in terms of the repair bandwidth. In Table 1, we compare the proposed construction with the previously known constructions.

The problem of constructing exact-repairable MDS codes with small repair bandwidth and small sub-packetization level has been previously addressed in [19, 27, 28]. We note that our construction shares some similarities with the constructions presented in [19, 28] as these constructions are obtained by introducing coupling among multiple independent codes as well. However, we work with the parity check matrix view (as opposed to the generator matrix view considered in [19, 28]) which ensures identical repair guarantees for all code blocks without distinguishing between systematic and parity blocks. For $r = n - k = 2$, Tamo and Efremenko construct exact-repairable MDS codes with near-optimal repair bandwidth and sub-packetization level $\ell = O(\log n)$ in [27]. The code construction obtained in [27] also satisfies the additional requirement that the same amount of data is downloaded from all the contacted blocks.

Exact repair of known codes with small repair bandwidth. The problem of devising exact repair mechanism with small repair bandwidth for known MDS codes has been studied in [10, 24, 30, 33]. In particular, [10, 24] consider the exact repair problem for the well-known Reed-Solomon codes. In [10], Guruswami and Wootters characterize optimal repair bandwidth for these codes in certain regimes of system parameters.

Locally repairable codes. Another line of work in distributed storage focuses on *locality*, the number of the code blocks contacted during the repair of a single code block, as a metric to characterize the efficiency of the repair process. The bounds on the failure tolerance of locally repairable codes, the codes with small locality, have been obtained in [7,12,16,20] and references therein. Furthermore, the constructions of locally repairable codes that are optimal with respect to these bounds are presented in [1, 6, 7, 12, 16, 20, 26]. Locally repairable codes that also minimize the repair bandwidth for repair of a code block are considered in [12, 20]. Here we note that the locally repairable codes are not MDS codes, and thus have extra storage overhead.

3 Exact-repairable MDS codes with near-optimal repair bandwidth

This paper aims to construct exact-repairable MDS vector codes with small sub-packetization level ℓ while achieving near-optimal repair bandwidth, i.e., incurring a small (multiplicative) loss in terms of the repair bandwidth. Towards this, we first introduce the notion of near-optimal repair bandwidth for MDS codes.

Definition 1. *Let $\mathcal{C}$ be an $[n, k\ell, d_{\min} = n - k + 1, \ell]_{|\mathbb{B}|}$ MDS vector code. We call $\mathcal{C}$ to be an (a, ℓ, d)-exact-repairable MDS code if for every $i \in [n]$ and $\mathbf{c} = (\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_n) \in \mathcal{C}$, we can perform exact repair of the code block $\mathbf{c}_i$ by contacting d other code blocks and downloading at most $a\big(\frac{d}{d-k+1}\big) \cdot \ell$ symbols (over $\mathbb{B}$) from the contacted code blocks.*

Remark 1. *It follows from the bound in* (5) *that for any exact-repairable MDS vector code we must have $a \geq 1$. Thus, $(a = 1, \ell, d)$-exact-repairable MDS codes correspond to the exact-repairable MDS vector codes with optimal repair bandwidth. Moreover, we say an MDS code has near-optimal repair bandwidth if it is an (a, ℓ, d)-exact-repairable MDS code for a small constant a.*

In this paper, we focus on the setting with $d = n - 1$, i.e., all the remaining $n - 1$ code blocks are contacted to repair a single code block. We now state the main result of this paper which summarizes the parameters of the codes constructed in this paper.

Theorem 1. *For an integer $1 \leq t \leq \lceil n/(n-k) \rceil - 1$ and a suitably chosen large enough field $\mathbb{B}$, the general construction presented in this paper gives $\big(1 + 1/t, (n-k)^t, n-1\big)$-exact-repairable MDS vector codes. Moreover, the obtained codes allow for repair-by-transfer schemes.*

We present our construction for $t = 1$, which gives $(2, n-k, n-1)$-exact-repairable MDS vector codes in Section 4. This construction conveys the main ideas behind our approach and establishes Theorem 1 for $t = 1$. The general construction which establishes Theorem 1 for all values of t is presented in Section 5.

Remark 2. *We note that for a given value of t, $\big(1 + 1/t\big)$ only serves as a clean upper bound on the repair bandwidth of the codes obtained in this paper. Specifically, if we substitute $t = \lceil n/(n-k) \rceil$ in general construction (cf. Section 5), we obtain $\big(1, (n-k)^{\lceil n/(n-k) \rceil}, n-1\big)$-exact-repairable MDS vector codes. This matches the best know sub-packetization level for optimal repair bandwidth, which is also near-optimal by* (7)*. In fact, in this case our construction specializes to the construction from [23].*

4 Construction of $(2, n-k, n-1)$-exact-repairable MDS code

In this section, we present a construction of exact-repairable MDS vector codes for all values of n and k. These codes have sub-packetization level $\ell = n-k$ and require $d = n-1$ code blocks during the repair process and have their repair bandwidth at most $2\big(\frac{n-1}{n-k}\big) \cdot \ell$, which is twice the cut-set bound (cf. (4)). The construction is described in Section 4.1. We illustrate the repair-by-transfer scheme for the obtained codes in Section 4.2. We argue the MDS property for the construction in Section 4.3.

4.1 Code construction

For an integer $a > 0$, we use $[a]$ to denote the set $\{1, 2, \ldots, a\}$. Let $r = n-k$. For ease of exposition, we assume that $r|n$ and $n = sr$. We partition the n code blocks in $r = n-k$ groups of size s each[2]. This partitioning allows us to index each code block by a tuple (u, v) where $u \in [r] = [n-k]$ and $v \in [s]$. In particular, for $i \in [n]$ the associated tuple (u, v) satisfies $i = (u-1)s + v$. With this notation in place, for $(u, v) \in [r] \times [s]$, we denote the $\big((u-1)s+v\big)$-th code block as

$$\mathbf{c}_{(u-1)s+v} = \mathbf{c}_{(u,v)} = \big(c(1;(u,v)), c(2;(u,v)), \ldots, c(r;(u,v))\big) \in \mathbb{B}^r.$$

Here, for $x \in [r]$, $c(x;(u,v))$ denotes the x-th symbol (over $\mathbb{B}$) of the $\big((u-1)s+v\big)$-th code block.

In order to construct an $[n, k\ell, d_{\min} = n-k+1, \ell = r = n-k]_{\mathbb{B}}$ MDS vector code $\mathcal{C}$, we specify an $r\ell \times n\ell$ (or $r^2 \times nr$ for our choice of ℓ) parity-check matrix $\mathbf{P}$ for the code $\mathcal{C}$. We classify the linear constraints defined by the parity-check matrix $\mathbf{P}$ into two types. Let $\{\lambda_i\}_{i\in[n]}$ be n distinct non-zero elements of $\mathbb{B}$.

- Type I **constraints:** We have r Type I constraints which are defined by the first r rows of the matrix $\mathbf{P}$. For every $x \in [r]$, we have

$$\sum_{(u,v)\in[r]\times[s]} c(x;(u,v)) = 0. \tag{8}$$

  In Example 1 below, the Type I constraints correspond to the identity blocks of the matrix $\mathbf{P}$ (cf. (11)).

- Type II **constraints:** Let ρ be an indeterminate which we specify later. We have $(r-1)\ell = (r-1)r$ Type II constraints which are defined as follows. For every $p \in \{1, \ldots, r-1\}$ and $x \in [r]$, we have

$$\underbrace{\sum_{(u,v)\in[r]\times[s]} \lambda^p_{(u-1)s+v} \cdot c(x;(u,v))}_{\text{(a)}} + \underbrace{\sum_{v\in[s]} \rho \cdot c(\overline{x+p};(x,v))}_{\text{(b)}} = 0, \tag{9}$$

  where for a strictly positive integer l, the quantity $\overline{l}$ is defined as follows.

$$\overline{l} = \begin{cases} r & \text{if } l \pmod r = 0, \\ l \pmod r & \text{otherwise.} \end{cases} \tag{10}$$

  We can partition the Type II constraints (cf. (9)) into $(r-1)$ groups (each group containing $\ell = r$ linear constraints) according to the value of $p \in \{1, \ldots, r-1\}$. In particular, r constraints associated with the same value of p constitute those r rows of the parity-check matrix $\mathbf{P}$ which are indexed by the set $\{pr+1, \ldots, (p+1)r\}$. (See the non-identity blocks of the matrix $\mathbf{P}$ in (11).)

[2]For a setting where $r \nmid n$, we can partition the n code blocks in $r = n-k$ groups, $n \pmod r$ groups with $\lceil \frac{n}{r} \rceil$ code blocks and the remaining groups with $\lfloor \frac{n}{r} \rfloor$ code blocks. The rest of the construction can be easily modified to work in this case as well.

Example 1. We illustrate the construction with an example. Assume that $n = 6$ and $k = 3$, i.e., $n - k = r = 3$. For these values of the system parameters, our 9×18 parity check matrix takes the following form.

$$\mathbf{P} = \left(\begin{array}{ccc|ccc|ccc|ccc|ccc|ccc} 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ \hline \lambda_1 & \rho & 0 & \lambda_2 & \rho & 0 & \lambda_3 & 0 & 0 & \lambda_4 & 0 & 0 & \lambda_5 & 0 & 0 & \lambda_6 & 0 & 0 \\ 0 & \lambda_1 & 0 & 0 & \lambda_2 & 0 & 0 & \lambda_3 & \rho & 0 & \lambda_4 & \rho & 0 & \lambda_5 & 0 & 0 & \lambda_6 & 0 \\ 0 & 0 & \lambda_1 & 0 & 0 & \lambda_2 & 0 & 0 & \lambda_3 & 0 & 0 & \lambda_4 & \rho & 0 & \lambda_5 & \rho & 0 & \lambda_6 \\ \hline \lambda_1^2 & 0 & \rho & \lambda_2^2 & 0 & \rho & \lambda_3^2 & 0 & 0 & \lambda_4^2 & 0 & 0 & \lambda_5^2 & 0 & 0 & \lambda_6^2 & 0 & 0 \\ 0 & \lambda_1^2 & 0 & 0 & \lambda_2^2 & 0 & \rho & \lambda_3^2 & 0 & \rho & \lambda_4^2 & 0 & 0 & \lambda_5^2 & 0 & 0 & \lambda_6^2 & 0 \\ 0 & 0 & \lambda_1^2 & 0 & 0 & \lambda_2^2 & 0 & 0 & \lambda_3^2 & 0 & 0 & \lambda_4^2 & 0 & \rho & \lambda_5^2 & 0 & \rho & \lambda_6^2 \end{array}\right). \quad (11)$$

The matrix $\mathbf{P}$ can be viewed as the perturbation of the block matrix $\mathbf{H}$ which is obtained by replacing all ρ entries in $\mathbf{P}$ with zeros. In particular, we can rewrite the matrix $\mathbf{P}$ as

$$\mathbf{P} = \mathbf{H} + \mathbf{E}^{\rho},$$

where $\mathbf{E}^{\rho}$ denotes the 9×18 matrix which contains all the ρ entries in $\mathbf{P}$ (cf. (11)) as its only non-zero entries. (See Figure 1.) Note that the block matrix $\mathbf{H}$ (with diagonal blocks) is a parity-check matrix of an $[n = 6, k\ell = 9, d_{\min} = 4, \ell = 3]_{\mathbb{B}}$ MDS vector code. Here, we also point out that the matrix $\mathbf{H}$ is defined by Type I constraints (cf. (8)) and the part (a) of the Type II constraints (cf. (9)). Similarly, the perturbation matrix $\mathbf{E}^{\rho}$ is defined by the part (b) of the Type II constraints (cf. (9)).

$$\mathbf{H} = \left(\begin{array}{ccc|ccc|ccc|ccc|ccc|ccc} 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ \hline \lambda_1 & 0 & 0 & \lambda_2 & 0 & 0 & \lambda_3 & 0 & 0 & \lambda_4 & 0 & 0 & \lambda_5 & 0 & 0 & \lambda_6 & 0 & 0 \\ 0 & \lambda_1 & 0 & 0 & \lambda_2 & 0 & 0 & \lambda_3 & 0 & 0 & \lambda_4 & 0 & 0 & \lambda_5 & 0 & 0 & \lambda_6 & 0 \\ 0 & 0 & \lambda_1 & 0 & 0 & \lambda_2 & 0 & 0 & \lambda_3 & 0 & 0 & \lambda_4 & 0 & 0 & \lambda_5 & 0 & 0 & \lambda_6 \\ \hline \lambda_1^2 & 0 & 0 & \lambda_2^2 & 0 & 0 & \lambda_3^2 & 0 & 0 & \lambda_4^2 & 0 & 0 & \lambda_5^2 & 0 & 0 & \lambda_6^2 & 0 & 0 \\ 0 & \lambda_1^2 & 0 & 0 & \lambda_2^2 & 0 & 0 & \lambda_3^2 & 0 & 0 & \lambda_4^2 & 0 & 0 & \lambda_5^2 & 0 & 0 & \lambda_6^2 & 0 \\ 0 & 0 & \lambda_1^2 & 0 & 0 & \lambda_2^2 & 0 & 0 & \lambda_3^2 & 0 & 0 & \lambda_4^2 & 0 & 0 & \lambda_5^2 & 0 & 0 & \lambda_6^2 \end{array}\right)$$

$$\mathbf{E}^{\rho} = \left(\begin{array}{ccc|ccc|ccc|ccc|ccc|ccc} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & \rho & 0 & 0 & \rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho & 0 & 0 & \rho & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho & 0 & 0 & \rho & 0 & 0 \\ \hline 0 & 0 & \rho & 0 & 0 & \rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \rho & 0 & 0 & \rho & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \rho & 0 & 0 & \rho & 0 \end{array}\right)$$

Figure 1: Illustration of matrices $\mathbf{H}$ and $\mathbf{E}^{\rho}$ in Example 1.

4.2 Exact repair of a code block

Let $(u^*, v^*) \in [r] \times [s]$ be the tuple associated with the code block to be repaired. Note that we need to reconstruct the r symbols $\big\{c(1;(u^*,v^*)), c(2;(u^*,v^*)), \ldots, c(r;(u^*,v^*))\big\}$. We divide the repair process in the following two stages.

1. First, we recover the symbol $c(u^*;(u^*,v^*))$ using the Type I constraint containing it (cf. (8)), i.e.,

$$\sum_{(u,v)\in[r]\times[s]} c(u^*;(u,v)) = 0. \tag{12}$$

   We download the $n-1$ symbols

$$\big\{c(u^*;(u,v)) \ : \ (u,v) \in [r]\times[s] \text{ s.t. } (u,v) \neq (u^*,v^*)\big\}$$

   from the remaining $d = n-1$ code blocks in this stage.

2. Next, we sequentially recover the $r-1$ symbols

$$\{c(1;(u^*,v^*)), \ldots, c(u^*-1;(u^*,v^*)), c(u^*+1;(u^*,v^*)), \ldots, c(r;(u^*,v^*))\}$$

   using the following $r-1$ Type II constraints (cf. (9)).

$$\underbrace{\sum_{(u,v)\in[r]\times[s]} \lambda^p_{(u-1)s+v} \cdot c(u^*;(u,v))}_{\text{(a)}} + \underbrace{\sum_{v\in[s]} \rho \cdot c(\overline{u^*+p};(u^*,v))}_{\text{(b)}} = 0 \ \text{ for } p \in \{1,\ldots,r-1\}. \tag{13}$$

   Note that the choice of Type I constraint used in the previous stage ensures that we now know all the values of the linear combinations in the part (a) of these Type II linear constraints. Now assuming that $p \in \{1,\ldots,r-1\}$ is such that $\overline{u^*+p} = \hat{u} \in [r]\backslash\{u^*\}$, by downloading the additional $s-1 = \frac{n}{r}-1$ symbols $\big\{c(\hat{u};(u^*,v) \ : \ v \in [s] \text{ s.t. } v \neq v^*\big\}$ which appear in the part (b) of the linear constraint associated with the underlying value of p, we can recover the desired symbol $c(\hat{u};(u^*,v^*))$. Thus, the entire second stage involves downloading the following number of symbols (in addition to the symbols downloaded in the first stage).

$$(r-1)(s-1) = (r-1)\,(n/r-1) \leq r\,(n/r-1) = n-r \leq n-1.$$

Note that the entire repair-by-transfer scheme described above downloads at most $2(n-1) = 2\left(\frac{n-1}{n-k}\right)\cdot\ell$ symbols (over $\mathbb{B}$), which is twice the cut-set bound (cf. (4)).

4.3 MDS property of the proposed codes

Next, we argue that the construction proposed in Section 4.1 gives us MDS vector codes. Let n and k be given system parameters. We show a way to choose the field $\mathbb{B}$ and assign a value to the indeterminate $\rho \in \mathbb{B}$ (cf. (9)) so that the parity-check matrix $\mathbf{P}$ defining the obtained code corresponds to a parity-check matrix of an $[n, k\ell = k(n-k), d_{\min} = n-k+1, \ell = n-k]_{\mathbb{B}}$ MDS vector code. This is equivalent to showing that for every $\mathcal{S} \subseteq [n]$ such that $|\mathcal{S}| = r = n-k$, the $r\ell \times r\ell$ sub-matrix $\mathbf{P}_{\mathcal{S}}$ (cf. (3)) is full rank. As illustrated in Example 1, the matrix $\mathbf{P}$ is obtained by perturbing a parity check matrix of an MDS vector code. In particular, we have

$$\mathbf{P} = \mathbf{H} + \mathbf{E}^{\rho}, \tag{14}$$

where $\mathbf{H}$ is a parity check matrix of an $[n, k\ell = k(n-k), d_{\min} = n-k+1, \ell = n-k]_{\mathbb{B}}$ MDS vector code. Assuming that we can find an irreducible polynomial of large enough degree over a field $\mathbb{L}$ with $|\mathbb{L}| \geq n+1$, it follows from Proposition 1 that one can select a field $\mathbb{B}$ and a non-zero element $\rho \in \mathbb{B}$ such that the matrix $\mathbf{P}$ corresponds to a parity-check matrix of an MDS vector code.

Proposition 1. *Let $\mathbb{L}$ be a field with at least $n+1$ elements and $\{\lambda_i\}_{i\in[n]}$ (cf. Section 4.1) be n distinct non-zero elements in the field $\mathbb{L}$. Assume that we can find an irreducible polynomial in $\mathbb{L}[X]$ of large enough degree. Then, one can construct a field $\mathbb{B}$ and select a non-zero element $\rho \in \mathbb{B}$ (cf.* (9)*) such that the matrix $\mathbf{P} = \mathbf{H} + \mathbf{E}^{\rho}$ (cf.* (14)*) is a parity-check matrix of an $[n, k\ell, d_{\min} = n-k+1, \ell]_{|\mathbb{B}|}$ MDS vector code.*

Proof. See Appendix B. □

The choice of ρ and $\mathbb{B}$ presented above gives us a fully explicit construction of the exact-repairable MDS vector codes (cf. Appendix B). However, this approach requires the size of the field $\mathbb{B}$ to be quite large. In particular, we need to have $|\mathbb{B}| \gg n^{(r-1)\ell}$. Next, we highlight another approach which ensures the existence of a choice for ρ such that the code obtained from the proposed construction (cf. Section 4.1) is an MDS vector code.

4.3.1 Random perturbations of a parity-check matrix of an MDS vector code

Note that the matrix $\mathbf{H}$ (cf. (14)) is a parity check matrix of an $[n, k\ell, d_{\min} = n-k+1, \ell = n-k]_{\mathbb{B}}$ code. If we randomly assign ρ to a non-zero element in $\mathbb{B}$, the parity check matrix $\mathbf{P}$ of the obtained code is a random perturbation of the matrix $\mathbf{H}$. Assuming that the field $\mathbb{B}$ has large enough size, it follows from Proposition 2 presented below that there exists a choice for the indeterminate ρ such that the matrix $\mathbf{P}$ corresponds to a parity-check matrix of an MDS vector code. This approach requires $|\mathbb{B}| \gg n^r r\ell$. We note that even though this alternative approach requires a slightly smaller field, it does not give us a fully explicit construction. Here, we also point out that Proposition 2 follows from the analysis presented in [23].

Proposition 2. *Assume that $\mathbb{B}$ is a field of large enough size and $\mathbf{H} \in \mathbb{B}^{r\ell \times n\ell}$ is a parity-check matrix of an $[n, k\ell, d_{\min} = n-k+1, \ell]_{|\mathbb{B}|}$ MDS vector code. Let $\mathbf{E}^{\rho}$ be a random $r\ell \times n\ell$ matrix with all of its non-zero entries equal to an element ρ which is selected uniformly at random from the non-zero elements in $\mathbb{B}$. Then, the probability that the matrix $\mathbf{P} = \mathbf{H} + \mathbf{E}^{\rho}$ is a parity-check matrix of an $[n, k\ell, d_{\min} = n-k+1, \ell]_{|\mathbb{B}|}$ MDS vector code is bounded away from zero.*

Proof. See Appendix C. □

5 Construction of $(1+\frac{1}{t}, (n-k)^t, n-1)$-exact-repairable MDS vector codes

In this section, we generalize the construction presented in Section 4. A design parameter t allows us to increase the sub-packetization level ℓ in order to decrease the repair bandwidth of the code. In particular, for the integer $1 \leq t \leq \lceil n/r \rceil - 1 = \lceil n/(n-k) \rceil - 1$, we design exact-repairable MDS vector codes with sub-packetization level $\ell = r^t = (n-k)^t$, $d = n-1$ and repair bandwidth at most

$$\left(1+\frac{1}{t}\right)\left(\frac{n-1}{n-k}\right)\cdot \ell \text{ symbols (over } \mathbb{B}).$$

This repair bandwidth is at most $\left(1+\frac{1}{t}\right)$ times the cut-set bound (cf. (4)). Recall that, for an integer $a > 0$, we use $[a]$ to denote the set $\{1, 2, \ldots, a\}$.

5.1 Code construction

Similar to Section 4.1, for ease of exposition, we assume that $r|n$ and $n = sr = s(n-k)$. We partition the n code blocks in $r = n-k$ groups of equal sizes with each group containing $s = \frac{n}{r} = \frac{n}{n-k}$ code blocks. Using this partition, we index each code block by a tuple (u, v) where $u \in [r] = [n-k]$ and $v \in [s]$. In particular, for $i \in [n]$ the associated tuple (u, v) satisfies

$$i = (u-1)s + v.$$

Furthermore, we index the $r^t = (n-k)^t$ symbols (over $\mathbb{B}$) in each code block by the r^t distinct t-length vectors in $[r]^t = [n-k]^t$. For $(u, v) \in [r] \times [s]$, the $\big((u-1)s+v\big)$-th code block can be represented as follows.

$$\mathbf{c}_{(u-1)s+v} = \mathbf{c}_{(u,v)} = \big\{c\big((x_1, x_2, \ldots, x_t); (u, v)\big)\big\}_{(x_1,x_2,\ldots,x_t)\in[r]^t}.$$

Let $\{\lambda_i\}_{i\in[n]}$ be n distinct non-zero elements of $\mathbb{B}$. We are now ready to present our construction of an $\big(1+\frac{1}{t}, \ell = r^t, d = n-1\big)$-exact-repairable MDS vector code $\mathcal{C}$ by defining an $r\ell \times n\ell$ parity-check matrix $\mathbf{P}$ of the code $\mathcal{C}$. Specifically, we classify the $r\ell = r^{t+1}$ linear constraints defined by the parity-check matrix $\mathbf{P}$ into two types.

- Type I **constraints:** We have $\ell = r^t$ Type I constraints which are defined by the first $\ell = r^t$ rows of the matrix $\mathbf{P}$. For every $(x_1, x_2, \ldots, x_t) \in [r]^t$, we have

$$\sum_{(u,v)\in[r]\times[s]} c\big((x_1, x_2, \ldots, x_t); (u, v)\big) = 0. \tag{15}$$

- Type II **constraints:** Let ρ be an indeterminate which we specify later. We have $(r-1)\ell = (r-1)r^t$ Type II constraints. Recall that for strictly positive integers l and m, the quantity $\overline{l}^{\{m\}}$ is defined as follows.

$$\overline{l}^{\{m\}} = \begin{cases} m & \text{if } l \pmod m = 0, \\ l \pmod m & \text{otherwise.} \end{cases} \tag{16}$$

  Assuming that $v \in [s] = \left[\frac{n}{r}\right]$ be such that $\overline{v}^{\{t\}} = a \in [t]$ and $p \in \{1, 2, \ldots, r-1\}$, we use $\overline{(x_1, x_2, \ldots, x_t)}^{v,p}$ to denote the vector obtained by modifying a single coordinate of the vector $(x_1, x_2, \ldots, x_t)$ in the following manner.

$$\begin{aligned} \overline{(x_1, x_2, \ldots, x_t)}^{v,p} &= (x_1, x_2, \ldots, x_{a-1}, \overline{x_{\overline{v}^{\{t\}}} + p}^{\{r\}}, x_{a+1}, \ldots, x_t) \\ &= (x_1, x_2, \ldots, x_{a-1}, \overline{x_a + p}^{\{r\}}, x_{a+1}, \ldots, x_t). \end{aligned}$$

  For every $p \in \{1, \ldots, r-1\}$ and $(x_1, x_2, \ldots, x_t) \in [r]^t$, we have an associated linear constraint in the parity-check matrix $\mathbf{P}$.

$$\underbrace{\sum_{(u,v)\in[r]\times[s]} \lambda^p_{(u-1)s+v} \cdot c\big((x_1, x_2, \ldots, x_t); (u, v)\big)}_{\text{(a)}}$$
$$+ \underbrace{\sum_{v\in[s]} \rho \cdot c\big(\overline{(x_1, x_2, \ldots, x_t)}^{v,p}; (x_{\overline{v}^{\{t\}}}, v)\big)}_{\text{(b)}} = 0. \tag{17}$$

We can partition the Type II constraints (cf. (17)) into $(r-1)$ groups (each group containing $\ell = r^t$ linear constraints) according to the value of $p \in \{1, \ldots, r-1\}$. In particular, r^t constraints associated with the same value of p constitute those r^t rows of the parity-check matrix $\mathbf{P}$ which are indexed by the set

$$\{pr^t + 1, \ldots, (p+1)r^t\} \subseteq [r\ell] = [r^{t+1}].$$

Example 2. In this example, we look at the composition of a Type II constraint (cf. 17) when $t = 2$. We assume that $n = 9$ and $r = n - k = 3$. This implies that $s = \frac{n}{n-k} = 3$. For $(x_1, x_2) \in [r]^t = [3]^2$ and $p = 1$ the associated Type II constraint takes the following form.

$$\underbrace{\sum_{(u,v)\in[3]\times[3]} \lambda_{(u-1)3+v} \cdot c\big((x_1,x_2);(u,v)\big)}_{\text{(a)}} \quad + \\ \underbrace{\rho \cdot c\big((\overline{x_1+1}^{\{3\}},x_2);(x_1,1)\big) + \rho \cdot c\big((x_1,\overline{x_2+1}^{\{3\}});(x_2,2)\big) + \rho \cdot c\big((\overline{x_1+1}^{\{3\}},x_2);(x_1,3)\big)}_{\text{(b)}} = 0. \tag{18}$$

Note that we have used the following equalities in (18) which hold for $t = 2$ and $s = \frac{n}{n-k} = 3$.

$$\overline{1}^{\{t=2\}} = \overline{3}^{\{2\}} = 1 \text{ and } \overline{2}^{\{2\}} = 2. \tag{19}$$

5.2 Exact repair of failed code blocks in the proposed codes

We now illustrate a mechanism to perform exact repair of a code block in the code obtained by the construction proposed in Section 5.1. Let $(u^*, v^*) \in [r] \times [s]$ be the tuple associated with the code block to be repaired. Note that we need to reconstruct the following r^t code symbols.

$$\big\{c\big((x_1, x_2, \ldots, x_t);(u^*, v^*)\big)\big\}_{(x_1,x_2,\ldots,x_t)\in[r]^t}. \tag{20}$$

Similar to Section 4.2, we divide the repair process in the following two stages.

1. In the first stage we utilize Type I constraints (cf. (15)) to recover the following r^{t-1} symbols.

$$\big\{c\big((x_1, \ldots, x_{a-1}, x_a = u^*, x_{a+1}, \ldots, x_t);(u^*, v^*)\big)\big\}_{(x_1,\ldots,x_{a-1},x_{a+1},\ldots,x_t)\in[r]^{t-1}}, \tag{21}$$

where $a = \overline{v^*}^{\{t\}}$. Recall that. for $(x_1, x_2, \ldots, x_t) \in [r]^t$, the Type I constraint takes the following form.

$$\sum_{(u,v)\in[r]\times[s]} c\big((x_1, x_2, \ldots, x_t));(u,v)\big) = 0. \tag{22}$$

Therefore, in order to recover the r^{t-1} symbols shown in (21) using these constraints, we download the following $(n-1)r^{t-1}$ symbols from the remaining $n-1$ code blocks.

$$\big\{c\left((x_1, \ldots, x_{a-1}, x_a = u^*, x_{a+1}, \ldots, x_t);(u,v)\right) \ : \ (x_1, \ldots, x_{a-1}, x_{a+1}, \ldots, x_t) \in [r]^{t-1} \\ \text{and } (u,v) \in [r]\times[s] \text{ s.t. } (u,v) \neq (u^*,v^*)\big\}, \tag{23}$$

where $a = \overline{v^*}^{\{t\}}$.

2. Note that, at the end of the stage 1 of the repair process, we have access to the following symbols (over $\mathbb{B}$) which also include the r^{t-1} symbols recovered in the stage 1.

$$\Big\{c\left((x_1,\ldots,x_{a-1},x_a=u^*,x_{a+1},\ldots,x_t);(u,v)\right) \;:\; (x_1,\ldots,x_{a-1},x_{a+1},\ldots,x_t)\in[r]^{t-1} \text{ and } (u,v)\in[r]\times[s]\Big\}. \quad (24)$$

In the stage 2 of the repair process, we employ the Type II constraints to sequentially recover the remaining $(r-1)r^{t-1}$ symbols

$$\Big\{c\big((x_1,\ldots,x_{a-1},x_a\neq u^*,x_{a+1},\ldots,x_t);(u^*,v^*)\big)\Big\}_{(x_1,\ldots,x_{a-1},x_{a+1},\ldots,x_t)\in[r]^{t-1}}, \quad (25)$$

where $a=\overline{v^*}^{\{t\}}$. Let $p\in\{1,2,\ldots,r-1\}$ be such that we have $\overline{u^*+p}^{\{r\}}=\hat{u}\in[r]\backslash\{u^*\}$. We utilize the following Type II constraint to repair the desired symbol $c\big((x_1,\ldots,x_{a-1},\hat{u},x_{a+1},\ldots,x_t);(u^*,v^*)\big)$.

$$\underbrace{\sum_{(u,v)\in[r]\times[s]}\lambda^p_{(u-1)s+v}\cdot c\big((x_1,\ldots,x_{a-1},x_a=u^*,x_{a+1},\ldots,x_t);(u,v)\big)+}_{\text{(a)}}$$
$$\underbrace{\rho\cdot c\big((x_1,\ldots,x_{a-1},\hat{u},x_{a+1},\ldots,x_t);(x_a=u^*,v^*)\big)}_{\text{(b-I)}}+$$
$$\underbrace{\sum_{v\in[s]\,:\,v\neq v^*}\rho\cdot c\big(\overline{(x_1,\ldots,x_a=u^*,\ldots,x_t)}^{v,p};(x_{\overline{v}^{\{t\}}},v)\big)}_{\text{(b-II)}}=0. \quad (26)$$

It is straightforward to verify that at the end of the stage 1 of the repair process, we know the value of the linear combination in the part (a) of this linear constraint (cf. (24)). We now argue that we also know many of the symbols appearing in the part (b-II) of this constraint. Note that the part (b-II) can be rewritten as follows.

$$\sum_{v\in[s]\,:\,v\neq v^*}\rho\cdot c\big(\overline{(x_1,\ldots,x_a=u^*,\ldots,x_t)}^{v,p};(x_{\overline{v}^{\{t\}}},v)\big)$$
$$=\underbrace{\sum_{v\neq v^*\,:\,\overline{v}^{\{t\}}=\overline{v^*}^{\{t\}}=a}\rho\cdot c\big((x_1,\ldots,x_{a-1},x_a=\hat{u},x_{a+1},\ldots,x_t);(x_a=u^*,v)\big)}_{\text{(b-II-1)}}$$
$$+\underbrace{\sum_{v\neq v^*\,:\,\overline{v}^{\{t\}}\neq\overline{v^*}^{\{t\}}=a}\rho\cdot c\big(\overline{(x_1,\ldots,x_a=u^*,\ldots,x_t)}^{v,p};(x_{\overline{v}^{\{t\}}},v)\big)}_{\text{(b-II-2)}}. \quad (27)$$

Note that the code symbols appearing in part (b-II-2) are indexed by the vectors which have their a-th coordinate equal to u^*. One can verify that these symbols are already known at the end of the stage 1 of the repair process (cf. (24)). Therefore, in order to recover the desired symbol

$$c\big((x_1,\ldots,x_{a-1},\hat{u},x_{a+1},\ldots,x_t);(u^*,v^*)\big)$$

using the linear constraint in (26), we need to only download the code symbols appearing in the part (b-II-1). Note that there are at most $\left\lfloor\frac{s}{t}\right\rfloor$ symbols in the part (b-II-1). Since we have to repair

$(r-1)r^{t-1}$ symbols in the stage 2 (cf. (25)), the number of symbols that we download in the stage 2 (in addition to the symbol downloaded in the stage 1) is at most

$$\begin{aligned}(r-1)r^{t-1}\left\lfloor\frac{s}{t}\right\rfloor &\leq (r-1)r^{t-1}\left(\frac{s}{t}\right)\\ &= \frac{r^{t-1}}{t}\frac{r-1}{r}n\\ &\overset{(i)}{\leq} \frac{r^{t-1}}{t}(n-1).\end{aligned} \tag{28}$$

Here the step (i) follows as, for $r = n-k \leq n$, we have $\frac{r-1}{r} \leq \frac{n-1}{n}$. Since we download $(n-1)r^{t-1}$ symbols during the stage 1 of the repair process, the total repair bandwidth is at most

$$(n-1)r^{t-1} + \frac{r^{t-1}}{t}(n-1) = \left(1+\frac{1}{t}\right)(n-1)r^{t-1} = \left(1+\frac{1}{t}\right)\left(\frac{n-1}{n-k}\right)\cdot\ell, \tag{29}$$

which is $(1+1/t)$ times the cut-set bound (cf. (4)).

5.3 MDS property of the proposed codes

The argument for this part is identical to that used in Section 4.3.

6 Conclusion and future directions

We construct MDS vector codes that allow for exact repair of a code block by downloading near-optimal amount of data from the remaining code blocks. These codes are well suited for distributed storage systems as they work with small sub-packetization level and enable repair-by-transfer mechanisms, where repair of a code block (node) requires minimal computation at the contacted code blocks (nodes). We conclude by pointing out a few directions to extend this work.

- **Reducing the size of base field $\mathbb{B}$.** The exact-repairability and the corresponding repair bandwidth of the proposed codes only depend on the combinatorial structure, i.e., the locations of non-zero entries, of the designed parity-check matrix. However, the argument which establishes the MDS property for these codes requires the size of the base field $\mathbb{B}$ to be quite large. We note that the similar issue arises in many previous works, e.g., [23, 28]. Recently, Ye and Barg have addressed this issue for the codes that operated exactly at the cut-set bound in [34, 35]. However, they again work with large sub-packetization level $n^{\left\lceil\frac{n}{n-k}\right\rceil}$. The reduction of the base field size for our construction is an interesting question, which has both theoretical and practical significance.

- **Constructing codes for general values of d.** In this paper we focus on the setting with $d = n-1$. Extending the construction proposed in this paper for general value of $k < d < n-1$ is another important direction to explore. Towards this, one relatively straightforward approach is to employ the ideas used in [22] to extend the construction from [23] to general values of d. For an integer $t \geq 1$, this would provide exact-repairable codes with sub-packetization level $(d-k+1)^t$ and small repair bandwidth. Moreover, the obtained codes would also have repair-by-transfer schemes.

- **Simultaneous repair of multiple code blocks.** The problem of designing MDS codes that allow for simultaneous repair of multiple code blocks has been addressed in several works, including [13, 21, 25, 34]. Designing codes that provide mechanisms to perform simultaneous repair of multiple code blocks, and as well as a good trade-off between the sub-packetization level and repair bandwidth is an interesting direction to pursue.

Acknowledgement

We would like to thank Itzhak Tamo for introducing us to the problem of constructing exact-repairable MDS codes with near-optimal repair bandwidth during the 2016 Information Theory and Applications (ITA) workshop. We are also grateful to him for commenting on an earlier version of this draft.

References

[1] M. Blaum, J. L. Hafner, and S. Hetzler. Partial-mds codes and their application to raid type of architectures. *IEEE Transactions on Information Theory*, 59(7):4510–4519, July 2013.

[2] V. R. Cadambe, C. Huang, J. Li, and S. Mehrotra. Polynomial length MDS codes with optimal repair in distributed storage. In *Proc. of Forty Fifth Asilomar Conference onSignals, Systems and Computers (ASILOMAR)*, pages 1850–1854, 2011.

[3] V. R. Cadambe, S. A. Jafar, H. Maleki, K. Ramchandran, and C. Suh. Asymptotic interference alignment for optimal repair of MDS codes in distributed storage. *IEEE Transactions on Information Theory*, 59(5):2974–2987, 2013.

[4] A. G. Dimakis, P. Godfrey, Y. Wu, M. Wainwright, and K. Ramchandran. Network coding for distributed storage systems. *IEEE Transactions on Information Theory*, 56(9):4539–4551, 2010.

[5] A. G. Dimakis, K. Ramchandran, Y. Wu, and C. Suh. A survey on network codes for distributed storage. *Proc. of the IEEE*, 99(3):476–489, 2011.

[6] P. Gopalan, C. Huang, B. Jenkins, and S. Yekhanin. Explicit maximally recoverable codes with locality. *IEEE Transactions on Information Theory*, 60(9):5245–5256, Sept 2014.

[7] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin. On the locality of codeword symbols. *IEEE Transactions on Information Theory*, 58(11):6925–6934, 2012.

[8] S. Goparaju, A. Fazeli, and A. Vardy. Minimum storage regenerating codes for all parameters. *CoRR*, abs/1602.04496, 2016.

[9] S. Goparaju, I. Tamo, and R. Calderbank. An improved sub-packetization bound for minimum storage regenerating codes. *IEEE Transactions on Information Theory*, 60(5):2770–2779, May 2014.

[10] V. Guruswami and M. Wootters. Repairing Reed-Solomon codes. *CoRR*, abs/1509.04764, 2015.

[11] W. Huang, M. Langberg, J. Kliewer, and J. Bruck. Communication efficient secret sharing. *CoRR*, abs/1505.07515, 2015.

[12] G. M. Kamath, N. Prakash, V. Lalitha, and P. V. Kumar. Codes with local regeneration and erasure correction. *IEEE Transactions on Information Theory*, 60(8):4637–4660, Aug 2014.

[13] A.-M. Kermarrec, N. Le Scouarnec, and G. Straub. Repairing multiple failures with coordinated and adaptive regenerating codes. In *Proceedings of 2011 International Symposium on Network Coding (NetCod)*, pages 1–6, 2011.

[14] R. Lidl and H. Niederreiter. *Introduction to Finite Fields and Their Applications*. Cambridge University Press, New York, NY, USA, 1986.

[15] F. J. MacWilliams and N. J. A. Sloane. *The Theory of Error-Correcting Codes*. Amsterdam: North-Holland, 1983.

[16] D. S. Papailiopoulos and A. G. Dimakis. Locally repairable codes. *IEEE Transactions on Information Theory*, 60(10):5843–5855, Oct 2014.

[17] D. S. Papailiopoulos, A. G. Dimakis, and V. Cadambe. Repair optimal erasure codes through hadamard designs. *IEEE Transactions on Information Theory*, 59(5):3021–3037, 2013.

[18] K. Rashmi, N. Shah, and P. Kumar. Optimal exact-regenerating codes for distributed storage at the MSR and MBR points via a product-matrix construction. *IEEE Transactions on Information Theory*, 57:5227–5239, 2011.

[19] K. V. Rashmi, N. B. Shah, and K. Ramchandran. A piggybacking design framework for read-and download-efficient distributed storage codes. In *Proc. of 2013 IEEE International Symposium on Information Theory (ISIT)*, pages 331–335, July 2013.

[20] A. S. Rawat, O. O. Koyluoglu, N. Silberstein, and S. Vishwanath. Optimal locally repairable and secure codes for distributed storage systems. *IEEE Transactions on Information Theory*, 60(1):212–236, 2014.

[21] A. S. Rawat, O. O. Koyluoglu, and S. Vishwanath. Centralized repair of multiple node failures with applications to communication efficient secret sharing. *CoRR*, abs/1603.04822, 2016.

[22] A. S. Rawat, O. O. Koyluoglu, and S. Vishwanath. Progress on high-rate MSR codes: Enabling arbitrary number of helper nodes. *CoRR*, abs/1601.06362, 2016.

[23] B. Sasidharan, G. K. Agarwal, and P. V. Kumar. A high-rate MSR code with polynomial sub-packetization level. *CoRR*, abs/1501.06662, 2015.

[24] K. Shanmugam, D. S. Papailiopoulos, A. G. Dimakis, and G. Caire. A repair framework for scalar MDS codes. *IEEE Journal on Selected Areas in Communications*, 32(5):998–1007, May 2014.

[25] K. W. Shum and Y. Hu. Cooperative regenerating codes. *IEEE Transactions on Information Theory*, 59(11):7229–7258, 2013.

[26] I. Tamo and A. Barg. A family of optimal locally recoverable codes. *IEEE Transactions on Information Theory*, 60(8):4661–4676, Aug 2014.

[27] I. Tamo and K. Efremenko. New results on msr codes. In *Information Theory and Applications Workshop (ITA), 2016*, Feb 2016.

[28] I. Tamo, Z. Wang, and J. Bruck. Zigzag codes: MDS array codes with optimal rebuilding. *IEEE Transactions on Information Theory*, 59(3):1597–1616, 2013.

[29] I. Tamo, Z. Wang, and J. Bruck. Access versus bandwidth in codes for storage. *IEEE Transactions on Information Theory*, 60(4):2028–2037, April 2014.

[30] Z. Wang, A. G. Dimakis, and J. Bruck. Rebuilding for array codes in distributed storage systems. In *2010 IEEE Globecom Workshops*, pages 1905–1909, Dec 2010.

[31] Z. Wang, I. Tamo, and J. Bruck. On codes for optimal rebuilding access. In *Proc. of the 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1374–1381, 2011.

[32] Z. Wang, I. Tamo, and J. Bruck. Long MDS codes for optimal repair bandwidth. In *Proc. of 2012 IEEE International Symposium on Information Theory (ISIT)*, pages 1182–1186, 2012.

[33] S. Xu, R. Li, P. P. C. Lee, Y. Zhu, L. Xiang, Y. Xu, and J. C. S. Lui. Single disk failure recovery for x-code-based parallel storage systems. *IEEE Transactions on Computers*, 63(4):995–1007, April 2014.

[34] M. Ye and A. Barg. Explicit constructions of high-rate MDS array codes with optimal repair bandwidth. *CoRR*, abs/1604.00454, 2016.

[35] M. Ye and A. Barg. Explicit constructions of optimal-access MDS codes with nearly optimal sub-packetization. *CoRR*, abs/1605.08630, 2016.

A Necessary sub-packetization level for MDS vector codes

Assume that $\mathbb{F} \cong \mathbb{B}^{\ell}$. Let $\mathcal{C} \subseteq \mathbb{F}^n$ be an MDS vector code with the sub-packetization level ℓ where all contacted nodes contribute to the repair process. For a constant $b \geq 1$, let the repair bandwidth of $\mathcal{C}$ for exact repair of a single code block is less than c times the cut-set bound, i.e.,

$$\text{No. of symbols (over } \mathbb{B}\text{) downloaded from the contacted } d = n-1 \text{ nodes} \leq b\left(\frac{n-1}{n-k}\right)\cdot \ell. \tag{30}$$

This implies that there exists at least one contacted node which contributes at most $\left\lfloor \frac{b\ell}{n-k} \right\rfloor$ symbols (over $\mathbb{B}$) during the repair process. Moreover, each of the contacted $d = n-1$ nodes sends at least 1 symbol (over $\mathbb{B}$) during the repair process. Hence, we have that

$$\left\lfloor \frac{b\ell}{n-k} \right\rfloor \geq 1. \tag{31}$$

This gives us that

$$\ell \geq \frac{n-k}{b} \text{ or } \ell = \Omega(n-k). \tag{32}$$

B Proof of Proposition 1.

Let $\mathbb{L}$ be a finite field such that all the n distinct non-zero elements $\{\lambda_i\}_{i\in[n]}$ used in the code construction (cf. Section 4.1) belong to $\mathbb{L}$. Furthermore, let $\rho \notin \mathbb{L}$ be an element from an extension of $\mathbb{L}$ such that its minimal polynomial $m_\rho(x) \in \mathbb{L}[X]$ has its degree $\deg(m_\rho)$ strictly greater than $(r-1)\ell$ [14]. We take $\mathbb{B}$ to be $\mathbb{L}(\rho)$, the simple extension of the field $\mathbb{L}$ to include ρ. Recall that $\mathbb{B} = \mathbb{L}(\rho) \cong \mathbb{L}[X]/\langle m_\rho(X)\rangle$, where $\langle m_\rho(X)\rangle \subset \mathbb{L}[X]$ is the ideal generated by the minimal polynomial $m_\rho(X)$. Moreover, $|\mathbb{B}| = |\mathbb{L}|^{\deg(m_\rho)}$.

Now, we argue that for such a choice of ρ and the associated field $\mathbb{B}$, the parity check matric $\mathbf{P}$ (cf. (14)) define an $[n, k\ell, d_{\min} = n-k+1, \ell]_{\mathbb{B}}$ MDS vector code, i.e., we have

$$\det(\mathbf{P}_{\mathcal{S}}) = \det(\mathbf{H}_{\mathcal{S}} + \mathbf{E}^{\rho}_{\mathcal{S}}) \neq 0 \ \forall \ \mathcal{S} \subseteq [n] \text{ such that } |\mathcal{S}| = r = n-k. \tag{33}$$

Recall that $\mathbf{H}$ is an $r\ell \times n\ell$ parity-check matrix of an $[n, k\ell, d_{\min} = n-k+1, \ell]_{|\mathbb{B}|}$ MDS vector code. Therefore, we have

$$\det\left(\mathbf{H}_{\mathcal{S}}\right) \neq 0 \quad \forall\, \mathcal{S} \subseteq [n] \text{ such that } |\mathcal{S}| = r. \tag{34}$$

Consider a set $\mathcal{S} \subseteq [n]$ such that $|\mathcal{S}| = r = n - k$ and the associated sub-matrix $\mathbf{H}_{\mathcal{S}} + \mathbf{E}_{\mathcal{S}}^{X}$ (cf. (14)), where X is an indeterminate. Note that the determinant of this $r\ell \times r\ell$ matrix can be expressed as

$$f_{\mathcal{S}}(X) = \det\left(\mathbf{H}_{\mathcal{S}} + \mathbf{E}_{\mathcal{S}}^{X}\right), \tag{35}$$

where $f_{\mathcal{S}}(X) \in \mathbb{L}[X]$ is a polynomial of degree at most $(r-1)\ell$ and its coefficients are defined by the elements $\{\lambda_i\}_{i\in\mathcal{S}} \subseteq \mathbb{L}$. We now argue that the polynomial $f_{\mathcal{S}}(X)$ is a non-trivial (not an identically zero) polynomial. Towards this, we consider the value of the polynomial $f_{\mathcal{S}}(X)$ at $X = 0$.

$$\begin{aligned} f_{\mathcal{S}}(0) &= \det\left(\mathbf{H}_{\mathcal{S}} + \mathbf{E}_{\mathcal{S}}^{0}\right) \\ &\overset{(i)}{=} \det\left(\mathbf{H}_{\mathcal{S}}\right) \\ &\overset{(ii)}{\neq} 0. \end{aligned} \tag{36}$$

Here the step (i) holds as $\mathbf{E}^0$ reduces to a zero matrix, and the step (ii) follows from (34). Since $f_{\mathcal{S}}(X)$ evaluates to a non-zero value at $X = 0$, it's a non-trivial polynomial. We now substitute $X = \rho$, which gives us the following (cf. 35).

$$\begin{aligned} f_{\mathcal{S}}(\rho) &= \det\left(\mathbf{H}_{\mathcal{S}} + \mathbf{E}_{\mathcal{S}}^{\rho}\right) \\ &\overset{(i)}{=} \det\left(\mathbf{P}_{\mathcal{S}}\right) \\ &\overset{(ii)}{\neq} 0. \end{aligned} \tag{37}$$

Here, the step (i) follows from the definition of $\mathbf{P}$ (cf. (14) and (33)). The step (ii) follows as we have that the degree of $f_{\mathcal{S}}(X) \in \mathbb{L}[X]$ is strictly less than the degree of $m_{\rho}(X) \in \mathbb{L}[X]$, the minimal polynomial of ρ. Since the choice of $\mathcal{S}$ is arbitrary over all the subsets of $[n]$ of size $r = n - k$. We have that

$$f_{\mathcal{S}}(\rho) = \det\left(\mathbf{P}_{\mathcal{S}}\right) \neq 0 \ \ \forall\, \mathcal{S} \subseteq [n] \text{ such that } |\mathcal{S}| = r. \tag{38}$$

This completes the proof. □

C Proof of Proposition 2.

Proof. Note that $\mathbf{H} \in \mathbb{B}^{r\ell \times n\ell}$ is an $r\ell \times n\ell$ parity-check matrix of an $[n, k\ell, d_{\min} = n - k + 1, \ell]_{|\mathbb{B}|}$ MDS vector code. Therefore, we have

$$\det\left(\mathbf{H}_{\mathcal{S}}\right) \neq 0 \quad \forall\, \mathcal{S} \subseteq [n] \text{ s.t. } |\mathcal{S}| = r. \tag{39}$$

Consider the perturbed matrix $\mathbf{P} = \mathbf{H} + \mathbf{E}^{\nu}$ where ν denotes an indeterminate. Let $\mathcal{S} \subseteq [n]$ such that $|\mathcal{S}| = r = n - k$ and $\mathbf{P}_{\mathcal{S}} = \mathbf{H}_{\mathcal{S}} + \mathbf{E}_{\mathcal{S}}^{\nu}$ be the associated sub-matrix (cf. (3)). Let $f_{\mathcal{S}}(\nu)$ be the determinant of the $r\ell \times r\ell$ matrix $\mathbf{P}_{\mathcal{S}}$, i.e.,

$$f_{\mathcal{S}}(\nu) = \det\left(\mathbf{P}_{\mathcal{S}}\right). \tag{40}$$

Note that $f_{\mathcal{S}}(\nu)$ is a polynomial in the indeterminate ν. Next, we argue that $f_{\mathcal{S}}(\nu)$ is a non-trivial (not an identically zero) polynomial. Let's consider the value of the polynomial $f_{\mathcal{S}}(\nu)$ at $\nu = 0$,

$$\begin{aligned} f_{\mathcal{S}}(0) &= \det\left(\mathbf{H}_{\mathcal{S}} + \mathbf{E}_{\mathcal{S}}^{\nu=0}\right) \\ &\overset{(i)}{=} \det\left(\mathbf{H}_{\mathcal{S}}\right) \end{aligned}$$

$$\stackrel{(ii)}{\neq} 0. \tag{41}$$

Here the step (i) holds as for $\nu = 0$, $\mathbf{E}^{\nu}$ reduces to a zero matrix. The step (ii) follows from (39). Since $f_{\mathcal{S}}(\nu)$ evaluates to a non-zero value at $\nu = 0$, it's a non-trivial polynomial. Note that this is true for any choice of the set $\mathcal{S} \subseteq [n]$ such that $|\mathcal{S}| = r = n - k$. We now consider the following polynomial.

$$g(\nu) = \prod_{\mathcal{S} \subseteq [n] \,:\, |\mathcal{S}| = r} f_{\mathcal{S}}(\nu). \tag{42}$$

Note that $g(\nu)$ is a non-trivial polynomial as it is a product of the non-trivial polynomials $\big\{f_{\mathcal{S}}(\nu)\big\}_{\mathcal{S} \subseteq [n] \,:\, |\mathcal{S}| = r}$. Moreover, we have that

$$\begin{aligned} \text{degree}\big(g(\nu)\big) &= \sum_{\mathcal{S} \subseteq [n] \,:\, |\mathcal{S}| = r} \text{degree}\big(f_{\mathcal{S}}(\nu)\big) \\ &\stackrel{(i)}{\leq} \binom{n}{r} r\ell, \end{aligned} \tag{43}$$

where (i) follows from the fact that for every subset $\mathcal{S}$ the degree of the associated polynomial $f_{\mathcal{S}}(\nu)$ is at most $r\ell$. Now, if we substitute ν with ρ which is selected uniformly at random from the non-zero elements in $\mathbb{B}^*$, then if follows from the Schwartz-Zippel lemma that we have

$$\begin{aligned} \mathbb{P}\big\{\mathbf{P} = \mathbf{H} + \mathbf{E}^{\rho} \text{ is not a parity-check matrix of an MDS vector code}\big\} &= \mathbb{P}\big\{g(a) = 0\big\} \\ &\leq \frac{\binom{n}{r} r\ell}{|\mathbb{B}| - 1}. \end{aligned} \tag{44}$$

Note that, for $|\mathbb{B}|$ large enough (in particular $|\mathbb{B}| = \Omega\left(\binom{n}{r} r\ell\right)$), the right hand side of (44) is strictly smaller than 1.

□